

Description

Method used to report a malfunction in a communication network.

The invention relates to a method for reporting a malfunction of network node in a communication network comprising a number of network nodes, as well as a network node for executing said inventive method.

Despite the use of high-quality components, the possibility cannot be completely excluded in a communication network of faults occurring in individual components or in the worst case, even of a total failure of the communication network occurring. In many cases the Simple Network Management Protocol, abbreviated to SNMP, is used to report faults. In accordance with the prior art there are two fault-reporting options.

Firstly network elements, especially network nodes, can be configured such that, with a malfunction a corresponding fault report is automatically created and transmitted to a higher-ranking network management station, where the fault is registered and further measures are deduced. For example routing tables can be adapted to the changed situation so that no data traffic is routed via faulty network elements. The fault message is sent in this case unsecured over the possibly faulty communication network. There is thus no guarantee of the fault report even actually reaching the intended network management stations.

The second option for fault or failure detection is the periodic interrogation of all network nodes by one or more network management stations. This method is also referred to by the term "polling". The disadvantage of this method is that a comparatively heavy load is imposed on the network management stations as well as on the communication network itself by

these ongoing interrogations. A further disadvantage is that a failure is detected on the basis of the lack of a response to a question. The lack of a response can namely also be because connection paths to a network node and not the network node itself are faulty. As long as the routing tables are not adapted accordingly in the communication network, incorrect fault reports can thus be produced.

The object of the invention is thus to specify a method in which a fault report securely reaches all network nodes of a communication network.

This occurs in accordance with the invention with a method of the type given at the start, in which each operable network node forwards a fault report of the faulty network node to all network nodes which are directly adjacent to the operable network node, except if a fault report is received which has already been forwarded.

These method steps are applied recursively for each network node, so that after a specific number of recursions, which is determined by the size of the network, all network nodes are informed about the fault. The method in accordance with the invention is thus heavily dependent on the Dijkstra shortest-path algorithm.

This algorithm is used to find the shortest connection between a start node and destination node, for route planners for example. In this case the entire network formed by nodes is investigated and thereby the shortest path is guaranteed to be found.

The fact that for fault reporting, a predetermined path, which under some circumstances can be faulty, is not used as in the prior art, but all network nodes are informed about a fault,

means that the fault report securely reaches its destination even with malfunctions of individual connections or network nodes. Only if a communication network is separated as a result of a fault in a number of subnetworks can a fault report not reach all network elements. However what was stated above continues to apply within a subnetwork.

It is advantageous for a faulty network node to send a fault report to all network nodes directly adjacent to it.

By contrast with the prior art, a fault report which is generated as a result of a fault recognized by a network node itself is not sent directly to a network management station but to all the network nodes adjacent to the faulty network node, which for their part forward this report. This is one way of reducing the load on the network management station.

- It is furthermore especially advantageous,
- if an operable network node continuously sends a message to be acknowledged to all the network nodes directly adjacent to it,
 - if this operable network node generates a fault report for a network node from which no acknowledgement was received, and
 - if this operable network node sends this fault report to all network nodes directly adjacent to it.

In this case the fault detection and fault reporting is relocated into the network nodes adjacent to a faulty network node. Each network node is in this case continuously interrogated by the network nodes adjacent to it as regards its operability. If this interrogation is not acknowledged as a result of a fault the interrogating network nodes assume that a fault has occurred. This method is also referred to by the term "polling". By contrast with the state-of-the-art, in which

polling is performed by a higher-ranking network management station, each network node is responsible here for its neighbors. Fault detection and reporting is thus undertaken locally.

- 5 A malfunction can advantageously be classified such that,
- the faulty report relates to the faulty network itself if all network nodes adjacent to it identify a fault, and
 - the fault report relates to the connection to the faulty network node if not all network nodes adjacent to it
- 10 identify a fault.

If the connection of a network node to all its neighboring network nodes fails it can be assumed to be highly probable that the network node itself is faulty or that the difference from the possibility that the network node is operating

15 correctly but all connections have simply failed is irrelevant for the routing. Under no circumstances should it be attempted to route data traffic via this network node. If however not all neighboring network nodes identify a fault only one or more connections to the network node have failed but the network

20 node itself is however operable. For this reason data traffic can continue to be routed via this network node, only the failed connections have to be taken into account.

A favorable variant of the invention is also produced by a method in which a network node which executes higher-ranking

25 functions deduces further steps from the reception of a fault report.

A fault report reaches all network nodes of a communication network and thereby also network nodes which are intended for higher-ranking functions. The latter can deduce further

30 measures from the reception of a fault report. Thus for example global routing tables can be adapted to the changed situation

so that no attempt is made to route data traffic via faulty network components.

The object to the invention is also achieved by a network node including

- 5 - means for receiving a fault report of a faulty network node,
- means for sending on this fault report to all network nodes directly adjacent to it,
- means for checking whether this fault report has already
10 been forwarded, and
- a transmission controller which is linked to the means the checking such that a fault report is only forwarded for a negative outcome of the check.

Provided each network node of a communication network or at
15 least sufficiently many are embodied in this way, a fault report reaches each point of the network. In this case only as much functionality is integrated into each network node as is absolutely necessary for forwarding a fault report. Since there must only be information available about directly adjacent
20 network nodes complex routing tables can be dispensed with. An inventive network node can thus be constructed in a comparatively simple technical manner.

It is advantageous in this case if the network node

- includes means for detecting a fault and
- 25 - includes means for sending a fault report to all network nodes directly adjacent to it.

As already mentioned a fault report, by contrast with the prior art, is not sent directly to a network management station but to all the network nodes directly adjacent to the faulty
30 network node. The interaction between the network nodes means that a fault report securely reaches each point of the

communication network.

It is also especially advantageous,

- if the network node includes means for continuously sending a message to be acknowledged to all the directly adjacent network nodes,
- if the network node includes means for generating a fault report for a network node from which no acknowledgement was received, and
- if the network node includes means for sending this fault report to all network nodes directly adjacent to it.

This type of network node it is suitable for monitoring all its neighbors. The polling is thus not performed by a specially designed network management station but decentrally by each network node.

- 15 An especially advantageous variant of the inventive network node is also produced,
- if this includes means for checking whether all network nodes adjacent to a faulty network node are identifying a fault, and
 - if this includes means the characterizing a fault report such that the fault report relates to the faulty network node itself if the result of the check is positive and otherwise the fault report relates to the connection to the faulty network node
- 25 As already mentioned, it can be assumed to be highly probable that a network node itself is faulty if the connection of this network node to all its neighboring network nodes fails. If however all neighboring network nodes identify a fault, then only one or more connections to the network node have failed,
- 30 the network node itself is however operable.

If is finally useful,

- if the network node includes means for executing higher-ranking functions, and
 - if the network node includes means for deducing further
- 5 steps from the reception of a fault report.

Higher-ranking functions can also be integrated into a network node. Thus global routing tables for example can be changed accordingly if a fault report is received. Network nodes embodied in this way are for example the network management

10 stations already mentioned.

At this point it should also be pointed out that the advantages stated for the inventive method are also applicable to the inventive network node

The invention will be explained in greater detail on the basis

15 of an exemplary embodiment shown in the Figures which relates to the detection of a fault as well as to the distribution of a corresponding fault report in a communication network.

The Figures show:

Figure 1: a communication network with a first through ninth

20 network node K1..K9;

Figure 2: the communication network with a faulty first network node K1 at a first point in time t1;

Figure 3: as Figure 2, but at a second point in time t2;

Figure 4: as Figure 2, but at a third point in time t3;

25 Figure 5: shows a communication network with "polling", with a fault in a connection between two network nodes;

Figure 6: as Figure 5, but with a fault in a first network node K1;

Figure 1 shows a communication network, comprising a first

30 through ninth network node K1..K9, with the first network node

K1 being connected to the second through fifth network node K2..K5, the second network node additionally being connected to the sixth and seventh network node K6 and K7, the third network node additionally being connected to the eighth and ninth network node K8 and K9 and finally the sixth network node being connected to the fifth and seventh network node K5 and K7.

The function of the arrangement shown in Figure 1 is now as follows:

Figure 2 shows the communication network at a first point in time t1. Assuming that a malfunction is established in the first network node K1, a corresponding fault report is sent to the directly adjacent network nodes, that is to the second through fifth network node K2..K5. The fault in this case is shown by a lightning symbol, the process of sending the fault report by arrow symbols. The sixth to ninth network node K6 to K9 have not yet received the fault report and are therefore shown by dotted outlines.

Figure 3 shows the communication network at a second point in time t2. In accordance with the inventive method the fault report is forwarded from each network node to the network nodes directly adjacent to it. This now means that the fault report is sent from the second network node K2 to the first, sixth and seventh network node K1, K6 and K7, from the third network node K3 to the first, eighth and ninth network node K1, K8 and K9, from the fourth network node K4 to the first network node K1 and finally from the fifth network node K5 to the first and sixth network node K1 and K6. In the mean time all network nodes K2..K9 have received the fault report and are thus shown by solid lines. The first network node K1 also receives a fault report with this method from the neighboring network nodes K2..K5, which it considers as acknowledgements to its own fault report.

Figure 4 shows the communication network at a third point in time t_3 . Here the fault report is sent in accordance with the methodology from the sixth network node K6 to the second, fifth and seventh network node K2, K5 and K7, from the seventh
5 network node K7 to the second and sixth network node K2 and K6, from the eighth network node K8 to the third network node K3 and finally from the ninth network node K9 to the third network node K3 as well. The second, third and fifth network node K2, K3 and K5 again consider the responses as acknowledgements to
10 their own fault report. The process for notifying the malfunction is completed at this point since each network node has forwarded the fault report to all network nodes adjacent to it, unless it has received the fault report at an earlier point in time. Therefore for example default report received from
15 that the sixth network node K6 by the seventh network node K7 is no longer forwarded since this has already been received from the second network node K2 at the point in time t_2 .

Figure 5 shows a communication network in which each network node K1..K9 is continuously sending messages to all neighboring
20 network nodes which confirm said messages, provided they are operable. If there is a fault this response naturally does not arrive. The method, which is also referred to as "polling", can be of advantage in such cases as when a network node for example is no longer in a position, as a result of total
25 failure, as shown in Figure 2, to issue a fault report. This "polling" is indicated by the corresponding arrows.

It is now assumed that there is a malfunction in the connection between the first and second network node K1 and K2. This is again indicated by a lightning symbol. The second network node
30 K2 thus does not receive any response from the first network node K1. Through communication with other network nodes K3, K5, which are adjacent to the first network node K1, it can

consequently be established that the connection from these nodes to the first network node K1 is functioning correctly. From this it can now be deduced that it is not the first network node K1 itself but only the connection between the
5 first and the second network node K1 and K2 which is faulty. The method steps mentioned can also be executed in a similar manner by the first network node in respect of the faulty connection.

Figure 6 likewise shows a communication network in which the
10 operability of the network nodes K1..K9 is monitored with the aid of polling. By contrast with Figure 5 however, none of the network nodes K2..K5 adjacent to the faulty network node K1 receive a response from this node. It can now be deduced from this that in this case a total failure of the first network
15 node K1 is involved.